

Abstract

The present invention provides a technique for skipping a locking process for an object in memory when a thread accesses an object that only it will access in order to reduce the load imposed on a system and to improve the overall system performance. A program executing in a computer system has multiple threads that share and access objects stored in memory. The objects have thread locality flags associated therewith that indicate the presence or absence of thread localities. The threads examine the thread locality flags for the objects they attempt to access to determine whether the corresponding objects, which are to be accessed, have localities for the threads. If, so the, threads skip the locking process and access objects immediately. If not, the object is locked prior to being accessed.